Interactive Wide-Area Visualization

Remote Visualization Console using Motion-JPEG Compression

Moving terascale datasets from supercomputing sites to remote users for visualization is extremely time-consuming. Therefore, it is best to leave data at the site where it is generated, and use visualization resources at that site for post-processing. However, images must be transferred from the visualization machine to the users at remote sites. When transferring such images over a wide area network (WAN), image quality, interactivity, and required WAN bandwidth must be carefully balanced. Although many mechanisms exist for transferring visualization displays to remote users, most of these methods use very large amounts of network bandwidth, or sacrifice interactivity to transfer images.

ASCI and DisCom have explored a number of techniques for implementing real-time, full frame-rate transfer of compressed images to remote visualization users. The technique described here uses commercial, off-the-shelf equipment to implement this capability. This demonstration uses visualization resources at Sandia, Los Alamos, and Lawrence Livermore Labs to run visualization codes that access locally-stored datasets, which are locally rendered and displayed real-time to the show floor. This system shows that full interactivity is maintained (at 30 frames per second), with little image feature degradation.

Figure 1: Original local console (left), and remote user's console (right)

The Problem

There is an imbalance between Storage Area Network (SAN) and Wide Area Network (WAN) bandwidth. In fact, SAN bandwidth often exceeds WAN bandwidth by at least two orders of magnitude. As a result, although terascale datasets can be quickly written from a supercomputer to storage, the amount of time required to transfer the simulation output to a remote user is considerably longer. This leads to a large amount of lag time between simulation completion and display of simulation of results to remote users.

For remote users, decreasing the lag between simulation completion and visualization requires the use of visualization resources that are local to the simulation host. For remote users to interact with the post-processed results, console images must be displayed across a wide-area network. However, for a 1280x1024 pixel resolution at 30 frames per second video stream, almost 1 Gigabit/second of wide area bandwidth is required. This exceeds the availability of reasonably-priced WAN bandwidth today.

In addition to the WAN bandwidth problem, remote users may not have high-performance visualization resources available on their sites. By providing such users with high-quality, remote access to visualization resources, they do not need to go to the expense of procuring duplicate resources (which may go underutilized) for their sites.

The Tradeoffs

When transferring video across a wide area network, three parameters must be carefully balanced to meet the end user's requirements – image resolution, interactivity, and WAN bandwidth. For a given value of WAN bandwidth, resolution and interactivity compete with each other.

Today's wide-area display solutions (e.g., raw X windows bitmaps, SGI's graphics rendering service, etc.) provide high resolution, but with reduced interactivity. This is due to the fact that WAN bandwidth simply cannot handle high resolution images at high framerates (interactivity). For some applications, this is fine. However, another option is needed.

To implement highly-interactive video distribution across relatively low-bandwidth (OC-3 and below) WAN connections, the resolution of the images must be reduced somewhat. A number of image compression algorithms exist that can handle this task efficiently. Furthermore, with the rise of networked video distribution, a number of products are now available to handle this

Sandia National Laboratories is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

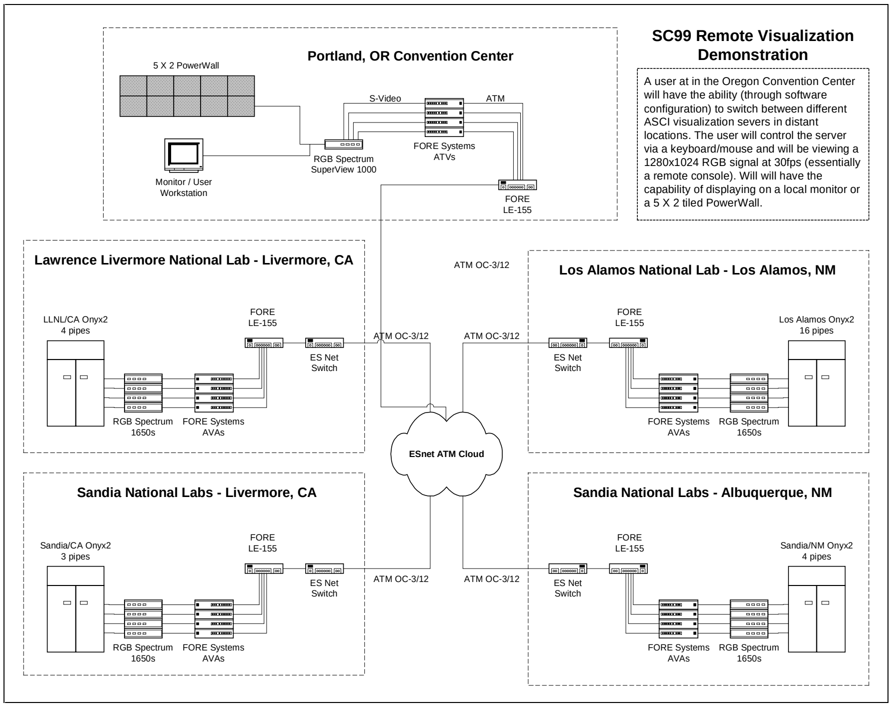

Figure 2: SC '99 demo configuration

over Asynchronous Transfer Mode (ATM) and Internet Protocol (IP) netowrks. However, these are lossy algorithms, so some image fidelity is sacrificed. However, with smart image compression algorithms, this sacrifice is barely noticeable to the end user.

Therefore, with compression, interactivity is preserved with little impact on image resolution.

Approach

The remote visualization system on display at SC '99 is shown in Figure 2. At each of the server sites, the 1280x1024 pixel frame buffers in each SGI Onyx 2 are converted to four 640x512 pixel RGB channels. Each of these RGB channels is connected to an RGB Spectrum 1650 scan converter, wich converts the RGB signal into an NTSC S-Video signal. Each of the S-Video signals is captured, compressed, and packed into ATM cells using FORE Systems AVA-300 video encoders.

The video streams are switched across the ESnet ATM network to the SC '99 SCInet network in Portland and to the ASCI booth. At this point, each of the four FORE Systems ATV-300 video decoders takes an ATM video stream, decompresses it, and displays it in NTSC S-Video format. The four S-Video signals are combined in an RGB Spectrum Superview 1000 multiple video windowing system, and displayed on a standard RGB monitor.

The user in the ASCI booth accesses the visualization machine using the standard IP remote access option (e.g., telnet, secure shell, etc.). In addition, keyboard and mouse inputs on the client side are sent to the visualization machine using standard X window protocols. However,instead of sending the visualization images through the X connection, they are converted, compressed, and transported across the ATM network.

Future Work

Although this is a useful first step in the development of compressed remote consoles, additional research is underway to correct some of the issues associated with this off-the-shelf implementation. These research initiatives include the following:

- Directly compressing RGB signals (without scan conversion to NTSC)
- Compressing digital video-formatted output
- Compressing and transmitting images between framebuffers without conversion to analog
- Using other image compression algorithms such as MPEG and JPEG-2000 (wavelet compression)

Summary

Terascale data sets from large simulation runs make data movement and processing cumbersome and expensive. Therefore, the best policy is to leave the data where it is generated, and use visualization resources at that site to post-process the data. However, this requires a mechanism to move visualization images back to the remote user.

Methods exist today to move these images, but they typically sacrifice real-time interaction to achieve high resolution. However, experience with compression algorithms shows that much image redundancy can be removed, and yet, retain many of the important features in the image. By compressing visualization images and transporting them across a network, commercial off-the-shelf devices can be used, and provide acceptable image resolution with no degradation in interactivity.

For more information contact: ….

Tom Tarman
Sandia National Laboratories
Albuquerque, New Mexico
tdtarma@sandia.gov

Tim Toole
Sandia National Laboratories
Livermore, California
ttoole@sandia.gov

Jerry Friesen
Sandia National Laboratories
Livermore, California
jafries@sandia.gov